

QUERY
 RESULTS

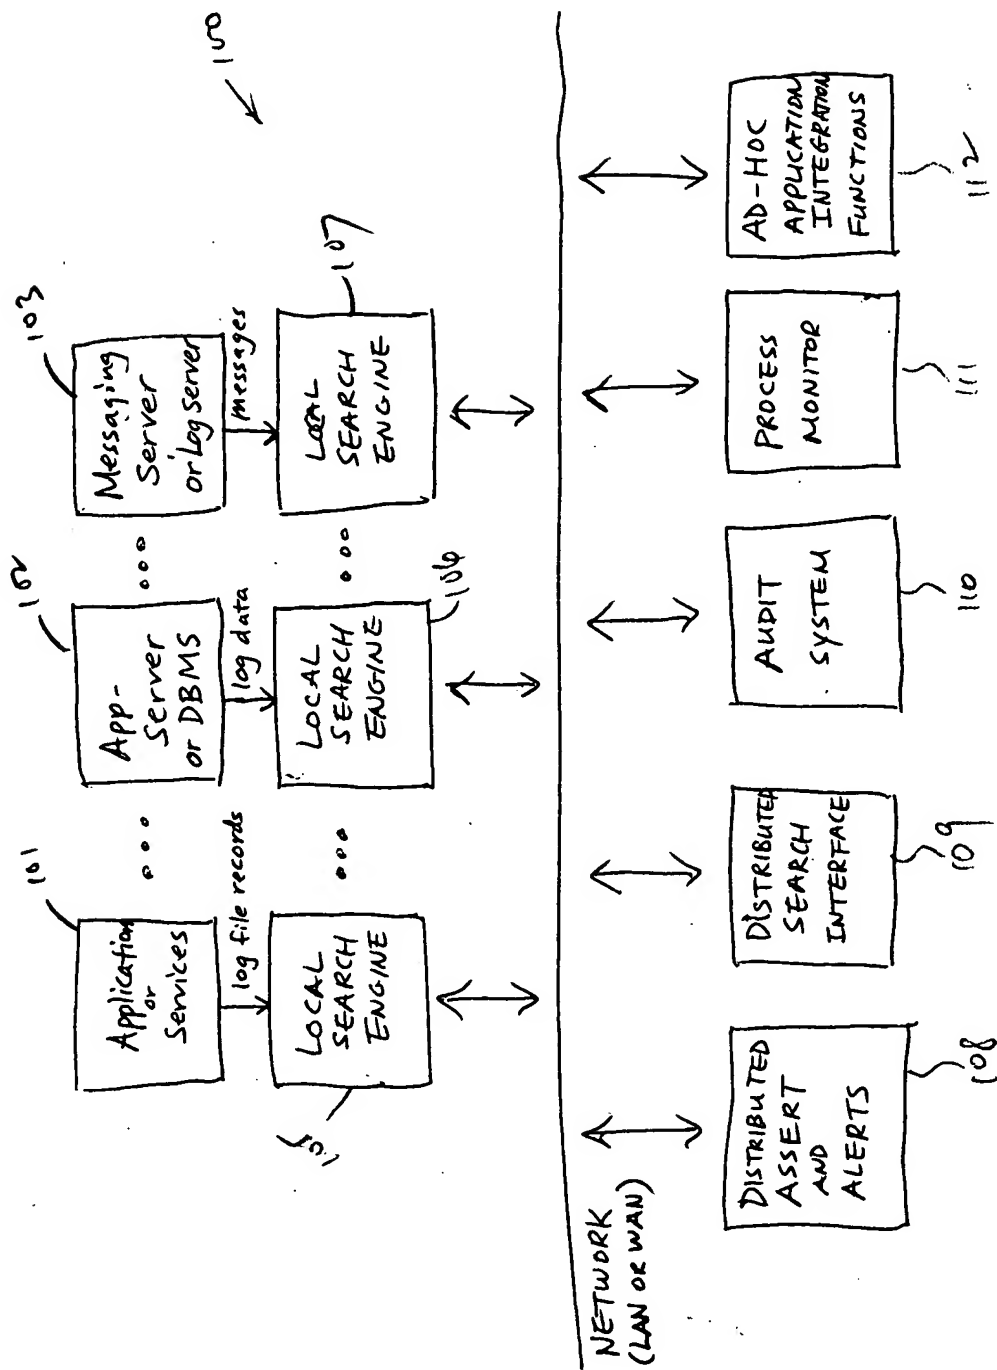


FIGURE 1

In-memory Index Format

Hashed by key

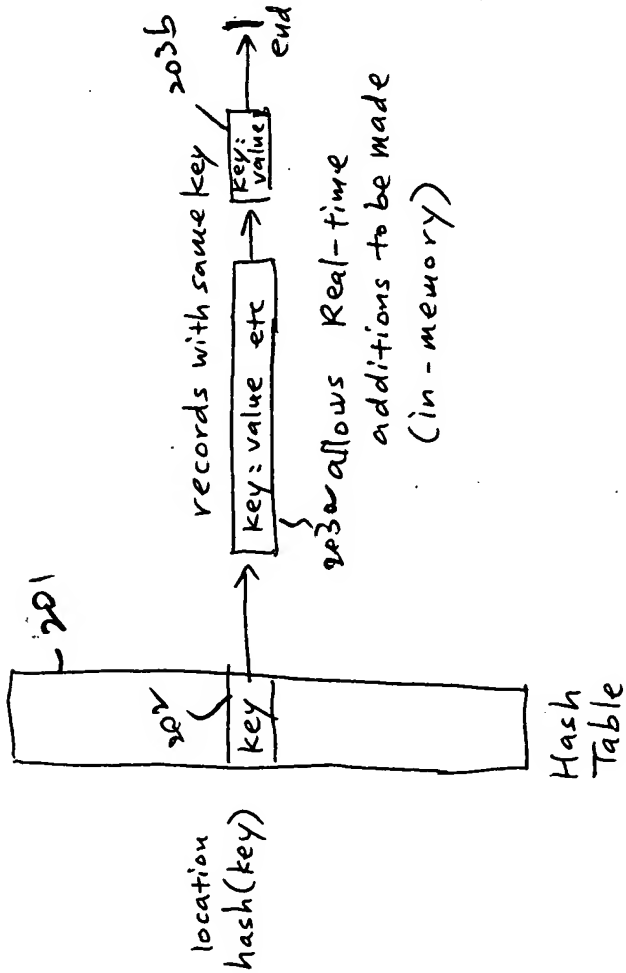


FIGURE 2

On-disk Index Format

Sorted by key




FIGURE 3

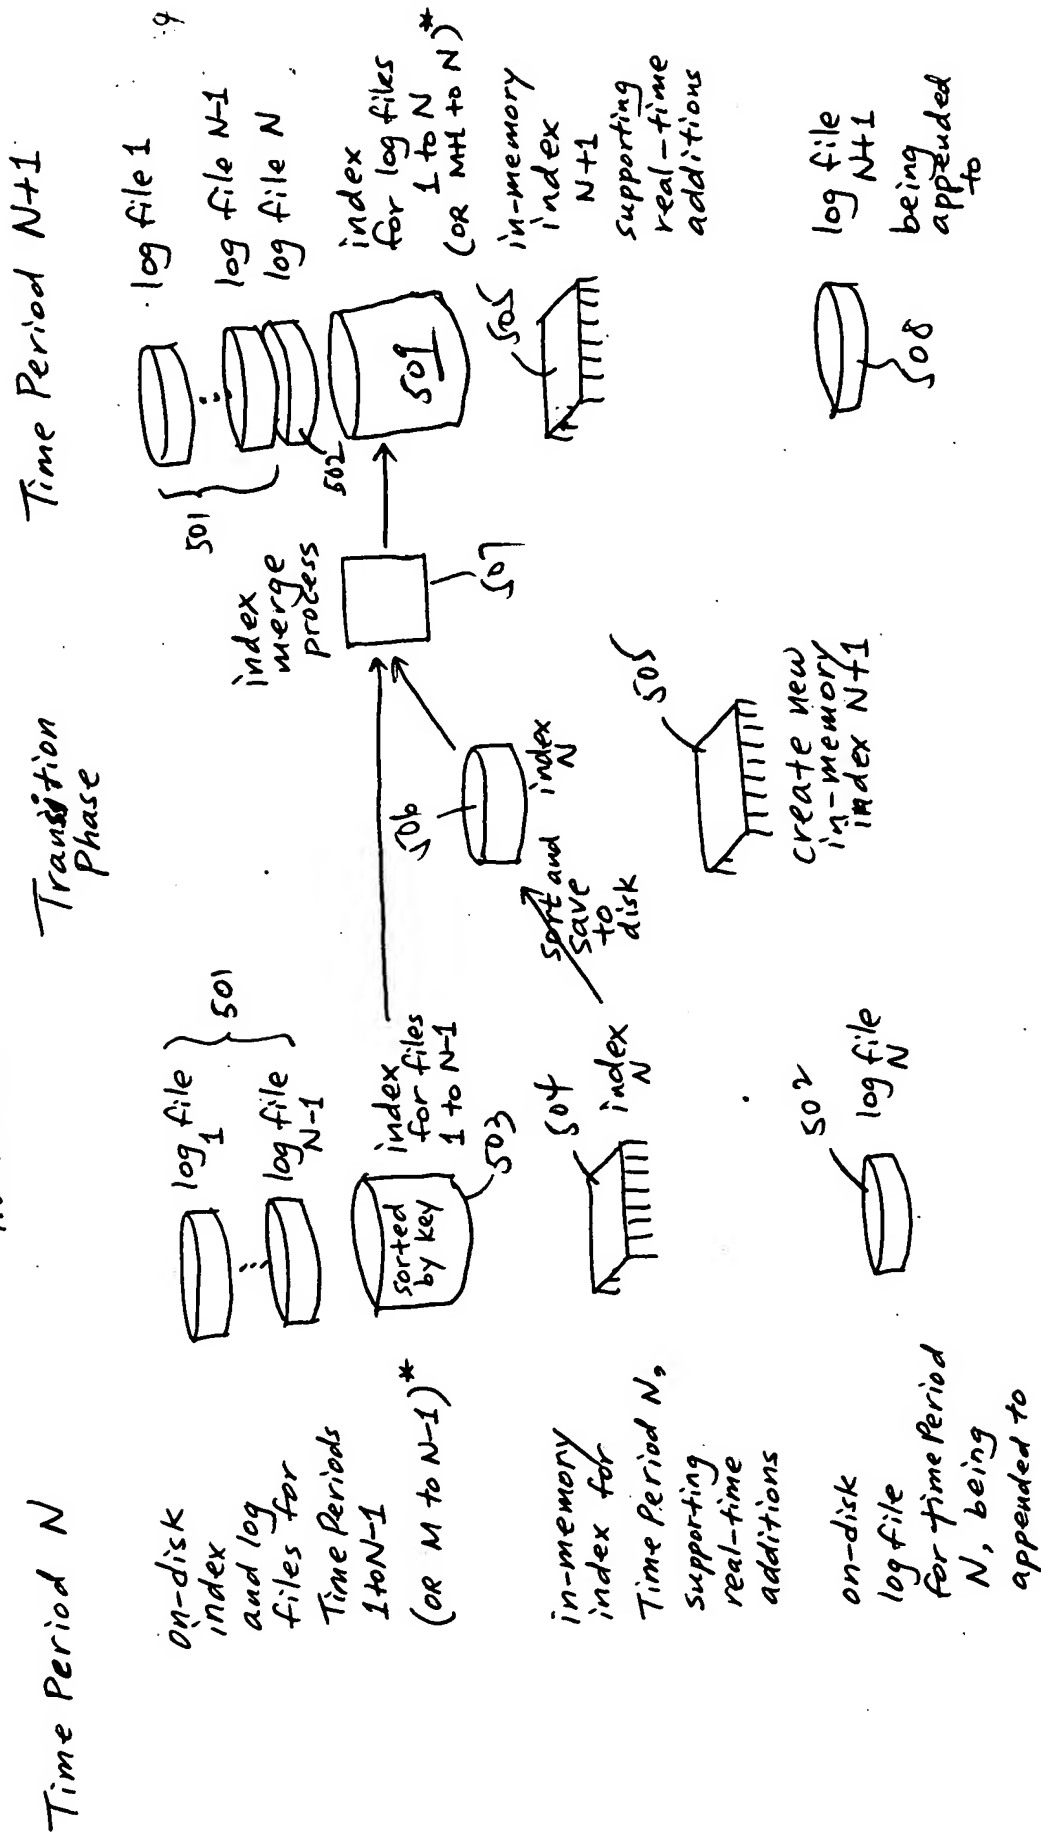
SAMPLE LOG DATA TO BE INDEXED

Log Entry ID: —
Time Stamp: —
Duration of computation: —
Module: —
Key Partial Results: —
Inputs: —
Outputs: —
Transaction ID: —
Customer ID: —
Keywords: —
Version: —
Side Effects: —
Server ID: —
Configuration: —
System Stats: —
Manager ID: —
Transaction Info: —

FIGURE 4

LOCAL SEARCH ENGINE

CONTINUOUS REAL-TIME INDEXING



BEST AVAILABLE COPY

FIGURE 5

DISTRIBUTED SEARCH INTERFACE

QUERY: CUSTOMER ID = X } 601
TIMESTAMP ≤ T₀
TIMESTAMP ≥ T₁
602

ASSERT: IF TransactionInfo/Paid THEN TransactionInfo/Delivered

RESULTS:

TransactionInfo	Customer ID	TimeStamp	Server	Module	LogEntryID
PAID	X	T ₁ ≤ T ≤ T ₀	host1	app1	13579
Delivered	X	.	host2	app2	24680
CREDIT CARD CHARGED	X
WAREHOUSE	X
UPS GROUND SENT	X
...	X
...	X

Automated Distributed Assert: if condition fail, alert operator/manager ID: 604

ANOTHER EXAMPLE QUERY: ZIP = 95070, Timestamp = T, 605
TransactionInfo/Price ≥ \$100
TransactionInfo / Call Center Report = URGENT-Missing

ANOTHER EXAMPLE: (Module = App1, Inputs = Y, KeyPartial Results/VariableName = Z) AND 606
(Module = App2, Inputs = W, KeyPartial Results/VariableName = Z)

FIGURE 6